

PROCESSING APPARATUS FOR DETERMINING WHICH PERSON IN A
GROUP IS SPEAKING

5 The present invention relates to the processing of image data to generate data to assist in archiving the image data.

10 The present invention further relates to the processing of image data and sound data to generate data to assist in archiving the image and sound data.

15 Many databases exist for the storage of data. However, the existing databases suffer from the problem that the ways in which the database can be interrogated to retrieve information therefrom are limited.

The present invention has been made with this problem in mind.

20 According to the present invention, there is provided an apparatus or method in which image and sound data recording the movements and speech of a number of people is processed using a combination of image processing and sound processing to identify which people shown in the
25 image data are speaking, and sound data is processed to

09532533-032200

[illegible]

10

15

20

25

The present invention further provides an apparatus or method for processing image data and sound data in such a system to identify a speaking participant.

5

The present invention further provides instructions, including in signal and recorded form, for configuring a programmable processing apparatus to become arranged as an apparatus, or to become operable to perform a method, in such a system.

10

According to the present invention, there is also provided an apparatus or method in which image data is processed to determine which person in the images is speaking by determining which person has the attention of the other people in the image, and sound data is processed to generate text data corresponding to the words spoken by the person using processing parameters selected in dependence upon the speaking participant identified by processing the image data.

15

20

The present invention also provides an apparatus or method in which image data is processed to determine at whom each person in the images is looking and to determine which of the people is speaking based thereon,

25

002220-032250

and sound data is processed to perform speech recognition for the speaking participant.

In this way, the speaking participant can be readily identified to enable the sound data to be processed.

The present invention further provides an apparatus or method for processing image data in such a system.

The present invention further provides instructions, including in signal and recorded form, for configuring a programmable processing apparatus to become arranged as an apparatus, or to become operable to perform a method, in such a system.

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 illustrates the recording of sound and video data from a meeting between a plurality of participants in a first embodiment;

Figure 2 is a block diagram showing an example of notional functional components within a processing

apparatus in the first embodiment;

Figure 3 shows the processing operations performed by processing apparatus 24 in Figure 2 prior to the meeting shown in Figure 1 between the participants starting;

Figure 4 schematically illustrates the data stored in meeting archive database 60 at step S2 and step S4 in Figure 3;

Figure 5 shows the processing operations performed at step S34 in Figure 3 and step S70 in Figure 7;

Figure 6 shows the processing operations performed at each of steps S42-1, S42-2 and S42-n in Figure 5;

Figure 7 shows the processing operations performed by processing apparatus 24 in Figure 2 while the meeting between the participants is taking place;

Figure 8 shows the processing operations performed at step S72 in Figure 7;

Figure 9 shows the processing operations performed at step S80 in Figure 8;

Figure 10 illustrates the viewing ray for a participant used in the processing performed at step S114 and step S124 in Figure 9;

5 Figure 11 illustrates the angles calculated in the processing performed at step S114 in Figure 9;

Figure 12 shows the processing operations performed at step S84 in Figure 8;

10 Figure 13 shows the processing operations performed at step S89 in Figure 8;

15 Figure 14 shows the processing operations performed at step S168 in Figure 13;

Figure 15 schematically illustrates the storage of information in the meeting archive database 60;

20 Figures 16A and 16B show examples of viewing histograms defined by data stored in the meeting archive database 60;

25 Figure 17 shows the processing operations performed at step S102 in Figure 8;

Year	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100
1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100	

10

15

20

25

25 processing apparatus 624 in Figure 22 prior to the

meeting shown in Figure 21 between the participants starting;

Figure 24 schematically illustrates the data stored in meeting archive database 660 at step S304 in Figure 23;

Figure 25 shows the processing operations performed at step S334 in Figure 23;

Figure 26 shows the processing operations performed by processing apparatus 624 in Figure 22 while the meeting between the participants is taking place;

Figure 27 shows the processing operations performed at step S372 in Figure 26;

Figure 28 shows the processing operations performed at step S380 in Figure 27;

Figure 29 illustrates the viewing ray for a participant used in the processing performed at step S414 in Figure 28;

Figure 30 illustrates the angles calculated in the processing performed at step S414 in Figure 28;

Figure 31 shows the processing operations performed at step S384 in Figure 27;

5 Figure 32 schematically illustrates the storage of information in the meeting archive database 660;

Figures 33A and 33B show examples of viewing histograms defined by data stored in the meeting archive database 660;

10 Figure 34 shows the processing operations performed by processing apparatus 624 to retrieve information from the meeting archive database 660;

15 Figure 35A shows the information displayed to a user at step S500 in Figure 34;

Figure 35B shows an example of information displayed to a user at step S504 in Figure 34; and

20 Figure 36 schematically illustrates a modification of the second embodiment in which a single database stores information from a plurality of meetings and is interrogated from one or more remote apparatus.

First Embodiment

Referring to Figure 1, a plurality of video cameras (three in the example shown in Figure 1, although this number may be different) 2-1, 2-2, 2-3 and a microphone array 4 are used to record image data and sound data respectively from a meeting taking place between a group of people 6, 8, 10, 12.

The microphone array 4 comprises an array of microphones arranged such that the direction of any incoming sound can be determined, for example as described in GB-A-2140558, US 4333170 and US 3392392.

The image data from the video cameras 2-1, 2-2, 2-3 and the sound data from the microphone array 4 is input via cables (not shown) to a computer 20 which processes the received data and stores data in a database to create an archive record of the meeting from which information can be subsequently retrieved.

Computer 20 comprises a conventional personal computer having a processing apparatus 24 containing, in a conventional manner, one or more processors, memory, sound card etc., together with a display device 26 and

user input devices, which, in this embodiment, comprise a keyboard 28 and a mouse 30.

The components of computer 20 and the input and output of data therefrom are schematically shown in Figure 2.

Referring to Figure 2, the processing apparatus 24 is programmed to operate in accordance with programming instructions input, for example, as data stored on a data storage medium, such as disk 32, and/or as a signal 34 input to the processing apparatus 24, for example from a remote database, by transmission over a communication network (not shown) such as the Internet or by transmission through the atmosphere, and/or entered by a user via a user input device such as keyboard 28 or other input device.

When programmed by the programming instructions, processing apparatus 24 effectively becomes configured into a number of functional units for performing processing operations. Examples of such functional units and their interconnections are shown in Figure 2. The illustrated units and interconnections in Figure 2 are, however, notional and are shown for illustration purposes only, to assist understanding; they do not necessarily

represent the exact units and connections into which the processor, memory etc of the processing apparatus 24 become configured.

5 Referring to the functional units shown in Figure 2, a central controller 36 processes inputs from the user input devices 28, 30 and receives data input to the processing apparatus 24 by a user as data stored on a storage device, such as disk 38, or as a signal 40
10 transmitted to the processing apparatus 24. The central controller 36 also provides control and processing for a number of the other functional units. Memory 42 is provided for use by central controller 36 and other functional units.

15 Head tracker 50 processes the image data received from video cameras 2-1, 2-2, 2-3 to track the position and orientation in three dimensions of the head of each of the participants 6, 8, 10, 12 in the meeting. In this
20 embodiment, to perform this tracking, head tracker 50 uses data defining a three-dimensional computer model of the head of each of the participants and data defining features thereof, which is stored in head model store 52, as will be described below.

Direction processor 53 processes sound data from the microphone array 4 to determine the direction or directions from which the sound recorded by the microphones was received. Such processing is performed in a conventional manner, for example as described in GB-A-2140558, US 4333170 and US 3392392.

Voice recognition processor 54 processes sound data received from microphone array 4 to generate text data therefrom. More particularly, voice recognition processor 54 operates in accordance with a conventional voice recognition program, such as "Dragon Dictate" or IBM "ViaVoice", to generate text data corresponding to the words spoken by the participants 6, 8, 10, 12. To perform the voice recognition processing, voice recognition processor 54 uses data defining the speech recognition parameters for each participant 6, 8, 10, 12, which is stored in speech recognition parameter store 56. More particularly, the data stored in speech recognition parameter store 56 comprises data defining the voice profile of each participant which is generated by training the voice recognition processor in a conventional manner. For example, the data comprises the data stored in the "user files" of Dragon Dictate after training.

Archive processor 58 generates data for storage in meeting archive database 60 using data received from head tracker 50, direction processor 53 and voice recognition processor 54. More particularly, as will be described below, video data from cameras 2-1, 2-2 and 2-3 and sound data from microphone array 4 is stored in meeting archive database 60 together with text data from voice recognition processor 54 and data defining at whom each participant in the meeting was looking at a given time.

Text searcher 62, in conjunction with central controller 36, is used to search the meeting archive database 60 to find and replay the sound and video data for one or more parts of the meeting which meet search criteria specified by a user, as will be described in further detail below.

Display processor 64 under control of central controller 36 displays information to a user via display device 26 and also replays sound and video data stored in meeting archive database 60.

Output processor 66 outputs part or all of the data from archive database 60, for example on a storage device such as disk 68 or as a signal 70.

Before beginning the meeting, it is necessary to initialise computer 20 by entering data which is necessary to enable processing apparatus 24 to perform the required processing operations.

5

Figure 3 shows the processing operations performed by processing apparatus 24 during this initialisation.

Referring to Figure 3, at step S1, central controller 36 causes display processor 64 to display a message on display device 26 requesting the user to input the names of each person who will participate in the meeting.

At step S2, upon receipt of data defining the names, for example input by the user using keyboard 28, central controller 36 allocates a unique identification number to each participant, and stores data, for example table 80 shown in Figure 4, defining the relationship between the identification numbers and the participants' names in the meeting archive database 60.

At step S3, central controller 36 causes display processor 64 to display a message on display device 26 requesting the user to input the name of each object at which a person may look for a significant amount of time

during the meeting, and for which it is desired to store archive data in the meeting archive database 60. Such objects may include, for example, a flip chart, such as the flip chart 14 shown in Figure 1, a whiteboard or blackboard, or a television, etc.

At step S4, upon receipt of data defining the names of the objects, for example input by the user using keyboard 28, central controller 36 allocates a unique identification number to each object, and stores data, for example as in table 80 shown in Figure 4, defining the relationship between the identification numbers and the names of the objects in the meeting archive database 60.

At step S6, central controller 36 searches the head model store 52 to determine whether data defining a head model is already stored for each participant in the meeting.

If it is determined at step S6 that a head model is not already stored for one or more of the participants, then, at step S8, central controller 36 causes display processor 64 to display a message on display device 26 requesting the user to input data defining a head model of each participant for whom a model is not already

stored.

In response, the user enters data, for example on a storage medium such as disk 38 or by downloading the data as a signal 40 from a connected processing apparatus, defining the required head models. Such head models may be generated in a conventional manner, for example as described in "An Analysis/Synthesis Cooperation for Head Tracking and Video Face Cloning" by Valente et al in Proceedings ECCV '98 Workshop on Perception of Human Action, University of Freiberg, Germany, June 6 1998.

At step S10, central controller 36 stores the data input by the user in head model store 52.

At step S12, central controller 36 and display processor 64 render each three-dimensional computer head model input by the user to display the model to the user on display device 26, together with a message requesting the user to identify at least seven features in each model.

In response, the user designates using mouse 30 points in each model which correspond to prominent features on the front, sides and, if possible, the back, of the participant's head, such as the corners of eyes,

nostrils, mouth, ears or features on glasses worn by the participant, etc.

At step S14, data defining the features identified by the user is stored by central controller 36 in head model store 52.

On the other hand, if it is determined at step S6 that a head model is already stored in head model store 52 for each participant, then steps S8 to S14 are omitted.

At step S16, central controller 36 searches speech recognition parameter store 56 to determine whether speech recognition parameters are already stored for each participant.

If it is determined at step S16 that speech recognition parameters are not available for all of the participants, then, at step S18, central controller 36 causes display processor 64 to display a message on display device 26 requesting the user to input the speech recognition parameters for each participant for whom the parameters are not already stored.

In response, the user enters data, for example on a

storage medium such as disk 38 or as a signal 40 from a remote processing apparatus, defining the necessary speech recognition parameters. As noted above, these parameters define a profile of the user's speech and are generated by training a voice recognition processor in a conventional manner. Thus for example, in the case of a voice recognition processor comprising Dragon Dictate, the speech recognition parameters input by the user correspond to the parameters stored in the "user files" of Dragon Dictate.

At step S20, data defining the speech recognition parameters input by the user is stored by central controller 36 in the speech recognition parameter store 56.

On the other hand, if it is determined at step S16 that the speech recognition parameters are already available for each of the participants, then steps S18 and S20 are omitted.

At step S22, central controller 36 causes display processor 64 to display a message on display device 26 requesting the user to perform steps to enable the cameras 2-1, 2-2 and 2-3 to be calibrated.

In response, the user carries out the necessary steps and, at step S24, central controller 36 performs processing to calibrate the cameras 2-1, 2-2 and 2-3. More particularly, in this embodiment, the steps performed by the user and the processing performed by central controller 36 are carried out in a manner such as that described in "Calibrating and 3D Modelling with a Multi-Camera System" by Wiles and Davison in 1999 IEEE Workshop on Multi-View Modelling and Analysis of Visual Scenes, ISBN 0769501109. This generates calibration data defining the position and orientation of each camera 2-1, 2-2 and 2-3 with respect to the meeting room and also the intrinsic parameters of each camera (aspect ratio, focal length, principal point, and first order radial distortion coefficient). The camera calibration data is stored, for example in memory 42.

At step S25, central controller 36 causes display processor 64 to display a message on display device 26 requesting the user to perform steps to enable the position and orientation of each of the objects for which identification data was stored at step S4 to be determined.

In response, the user carries out the necessary steps

and, at step S26, central controller 36 performs processing to determine the position and orientation of each object. More particularly, in this embodiment, the user places coloured markers at points on the perimeter of the surface(s) of the object at which the participants in the meeting may look, for example the plane of the sheets of paper of flip chart 14. Image data recorded by each of cameras 2-1, 2-2 and 2-3 is then processed by central controller 36 using the camera calibration data stored at step S24 to determine, in a conventional manner, the position in three-dimensions of each of the coloured markers. This processing is performed for each camera 2-1, 2-2 and 2-3 to give separate estimates of the position of each coloured marker, and an average is then determined for the position of each marker from the positions calculated using data from each camera 2-1, 2-2 and 2-3. Using the average position of each marker, central controller 36 calculates in a conventional manner the centre of the object surface and a surface normal to define the orientation of the object surface. The determined position and orientation for each object is stored as object calibration data, for example in memory 42.

At step S27, central controller 36 causes display

processor 64 to display a message on display device 26 requesting the next participant in the meeting (this being the first participant the first time step S27 is performed) to sit down.

5

At step S28, processing apparatus 24 waits for a predetermined period of time to give the requested participant time to sit down, and then, at step S30, central controller 36 processes the respective image data from each camera 2-1, 2-2 and 2-3 to determine an estimate of the position of the seated participant's head for each camera. More particularly, in this embodiment, central controller 36 carries out processing separately for each camera in a conventional manner to identify each portion in a frame of image data from the camera which has a colour corresponding to the colour of the skin of the participant (this colour being determined from the data defining the head model of the participant stored in head model store 52), and then selects the portion which corresponds to the highest position in the meeting room (since it is assumed that the head will be the highest skin-coloured part of the body). Using the position of the identified portion in the image and the camera calibration parameters determined at step S24, central controller 36 then determines an estimate of the three-

10

15

20

25

dimensional position of the head in a conventional manner. This processing is performed for each camera 2-1, 2-2 and 2-3 to give a separate head position estimate for each camera.

5

At step S32, central controller 36 determines an estimate of the orientation of the participant's head in three dimensions for each camera 2-1, 2-2 and 2-3. More particularly, in this embodiment, central controller 36 renders the three-dimensional computer model of the participant's head stored in head model store 52 for a plurality of different orientations of the model to produce a respective two-dimensional image of the model for each orientation. In this embodiment, the computer model of the participant's head is rendered in 108 different orientations to produce 108 respective two-dimensional images, the orientations corresponding to 36 rotations of the head model in 10° steps for each of three head inclinations corresponding to 0° (looking straight ahead), $+45^\circ$ (looking up) and -45° (looking down). Each two-dimensional image of the model is then compared by central processor 36 with the part of the video frame from a camera 2-1, 2-2, 2-3 which shows the participant's head, and the orientation for which the image of the model best matches the video image data is

10

15

20

25

002250" E22550

selected, this comparison and selection being performed for each camera to give a head orientation estimate for each camera. When comparing the image data produced by rendering the head model with the video data from a camera, a conventional technique is used, for example as described in "Head Tracking Using a Textured Polygonal Model" by Schödl, Haro & Essa in Proceedings 1998 Workshop on Perceptual User Interfaces.

At step S34, the respective estimates of the position of the participant's head generated at step S30 and the respective estimates of the orientation of the participant's head generated at step S32 are input to head tracker 50 and frames of image data received from each of cameras 2-1, 2-2 and 2-3 are processed to track the head of the participant. More particularly, in this embodiment, head tracker 50 performs processing to track the head in a conventional manner, for example as described in "An Analysis/Synthesis Cooperation for Head Tracking and Video Face Cloning" by Valente et al in Proceedings EECV '98 Workshop on Perception of Human Action, University of Freiberg, Germany, June 6 1998.

Figure 5 summarises the processing operations performed by head tracker 50 at step S34.

Referring to Figure 5, in each of steps S42-1 to S42-n ("n" being three in this embodiment since there are three cameras), head tracker 50 processes image data from a respective one of the cameras recording the meeting to determine the positions of the head features of the participant (stored at step S14) in the image data from the camera and to determine therefrom the three-dimensional position and orientation of the participant's head for the current frame of image data from that camera.

Figure 6 shows the processing operations performed at a given one of steps S42-1 to S42-n, the processing operations being the same at each step but being carried out on image data from a different camera.

Referring to Figure 6, at step S50, head tracker 50 reads the current estimates of the 3D position and orientation of the participant's head, these being the estimates produced at steps S30 and S32 in Figure 3 the first time step S50 is performed.

At step S52, head tracker 50 uses the camera calibration data generated at step S24 to render the three-dimensional computer model of the participant's head

stored in head model store 52 in accordance with the estimates of position and orientation read at step S50.

At step S54, head tracker 50 processes the image data for the current frame of video data received from the camera to extract the image data from each area which surrounds the expected position of one of the head features identified by the user and stored at step S14, the expected positions being determined from the estimates read at step S50 and the camera calibration data generated at step S24.

At step S56, head tracker 50 matches the rendered image data generated at step S52 and the camera image data extracted at step S54 to find the camera image data which best matches the rendered head model.

At step S58, head tracker 50 uses the camera image data identified at step S56 which best matches the rendered head model together with the camera calibration data stored at step S24 (Figure 3) to determine the 3D position and orientation of the participant's head for the current frame of video data.

Referring again to Figure 5, at step S44, head tracker 50

uses the camera image data identified at each of steps S42-1 to S42-n which best matches the rendered head model (identified at step S58 in Figure 6) to determine an average 3D position and orientation of the participant's head for the current frame of video data.

At the same time that step S44 is performed, at step S46, the positions of the head features in the camera image data determined at each of steps S42-1 to S42-n (identified at step S58 in Figure 6) are input into a conventional Kalman filter to generate an estimate of the 3D position and orientation of the participant's head for the next frame of video data. Steps S42 to S46 are performed repeatedly for the participant as frames of video data are received from video camera 2-1, 2-2 and 2-3.

Referring again to Figure 3, at step S36, central controller 36 determines whether there is another participant in the meeting, and steps S27 to S36 are repeated until processing has been performed for each participant in the manner described above. However, while these steps are performed for each participant, at step S34, head tracker 50 continues to track the head of each participant who has already sat down.

When it is determined at step S36 that there are no further participants in the meeting and that accordingly the head of each participant is being tracked by head tracker 50, then, at step S38, central controller 36 causes an audible signal to be output from processing apparatus 24 to indicate that the meeting between the participants can begin.

Figure 7 shows the processing operations performed by processing apparatus 24 as the meeting between the participants takes place.

Referring to Figure 7, at step S70, head tracker 50 continues to track the head of each participant in the meeting. The processing performed by head tracker 50 at step S70 is the same as that described above with respect to step S34, and accordingly will not be described again here.

At the same time that head tracker 50 is tracking the head of each participant at step S70, at step S72 processing is performed to generate and store data in meeting archive database 60.

Figure 8 shows the processing operations performed at

step S72.

Referring to Figure 8, at step S80, archive processor 58 generates a so-called "viewing parameter" for each participant defining at which person or which object the participant is looking.

Figure 9 shows the processing operations performed at step S80.

Referring to Figure 9, at step S110, archive processor 58 reads the current three-dimensional position of each participant's head from head tracker 50, this being the average position generated in the processing performed by head tracker 50 at step S44 (Figure 5).

At step S112, archive processor 58 reads the current orientation of the head of the next participant (this being the first participant the first time step S112 is performed) from head tracker 50. The orientation read at step S112 is the average orientation generated in the processing performed by head tracker 50 at step S44 (Figure 5).

At step S114, archive processor 58 determines the angle

between a ray defining where the participant is looking (a so-called "viewing ray") and each notional line which connects the head of the participant with the centre of the head of another participant.

5

More particularly, referring to Figures 10 and 11, an example of the processing performed at step S114 is illustrated for one of the participants, namely participant 6 in Figure 1. Referring to Figure 10, the orientation of the participant's head read at step S112 defines a viewing ray 90 from a point between the centre of the participant's eyes which is perpendicular to the participant's head. Similarly, referring to Figure 11, the positions of all of the participant's heads read at step S110 define notional lines 92, 94, 96 from the point between the centre of the eyes of participant 6 to the centre of the heads of each of the other participants 8, 10, 12. In the processing performed at step S114, archive processor 58 determines the angles 98, 100, 102 between the viewing ray 90 and each of the notional lines 92, 94, 96.

10

15

20

25

Referring again to Figure 9, at step S116, archive processor 58 selects the angle 98, 100 or 102 which has the smallest value. Thus, referring to the example shown

in Figure 11, the angle 100 would be selected.

At step S118, archive processor 58 determines whether the angle selected at step S116 has a value less than 10° .

5

If it is determined at step S118 that the angle is less than 10° , then, at step S120, archive processor 58 sets the viewing parameter for the participant to the identification number (allocated at step S2 in Figure 3) of the participant connected by the notional line which makes the smallest angle with the viewing ray. Thus, referring to the example shown in Figure 11, if angle 100 is less than 10° , then the viewing parameter would be set to the identification number of participant 10 since angle 100 is the angle between viewing ray 90 and notional line 94 which connects participant 6 to participant 10.

10

15

20

On the other hand, if it is determined at step S118 that the smallest angle is not less than 10° , then, at step S122, archive processor 58 reads the position of each object previously stored at step S26 (Figure 3).

25

At step S124, archive processor 58 determines whether the viewing ray 90 of the participant intersects the plane of

any of the objects.

If it is determined at step S124 that the viewing ray 90 does intersect the plane of an object, then, at step S126, archive processor 50 sets the viewing parameter for the participant to the identification number (allocated at step S4 in Figure 3) of the object which is intersected by the viewing ray, this being the nearest intersected object to the participant if more than one object is intersected by the viewing ray 90.

On the other hand, if it is determined at step S124 that the viewing ray 90 does not intersect the plane of an object, then, at step S128, archive processor 58 sets the value of the viewing parameter for the participant to "0". This indicates that the participant is determined to be looking at none of the other participants (since the viewing ray 90 is not close enough to any of the notional lines 92, 94, 96) and none of the objects (since the viewing ray 90 does not intersect an object). Such a situation could arise, for example, if the participant was looking at some object in the meeting room for which data had not been stored at step S4 and which had not been calibrated at step S26 (for example the notes held by participant 12 in the example shown in Figure 1).

At step S130, archive processor 58 determines whether there is another participant in the meeting, and steps S112 to S130 are repeated until the processing described above has been carried out for each of the participants.

5

Referring again to Figure 8, at step S82, central controller 36 and voice recognition processor 54 determine whether any speech data has been received from the microphone array 4 corresponding to the current frame of video data.

10

If it is determined at step S82 that speech data has been received, then, at step S84, processing is performed to determine which of the participants in the meeting is speaking.

15

Figure 12 shows the processing operations performed at step S84.

20

Referring to Figure 12, at step S140, direction processor 53 processes the sound data from the microphone array 4 to determine the direction or directions from which the speech is coming. This processing is performed in a conventional manner, for example as described in GB-A-2140558, US 4333170 and US 3392392.

25

At step S142, archive processor 58 reads the position of each participant's head determined by head tracker 50 at step S44 (Figure 5) for the current frame of image data and determines therefrom which of the participants has a head at a position corresponding to a direction determined at step S140, that is, a direction from which the speech is coming.

At step S144, archive processor 58 determines whether there is more than one participant in a direction from which the speech is coming.

If it is determined at step S144 that there is only one participant in the direction from which the speech is coming, then, at step S146, archive processor 58 selects the participant in the direction from which the speech is coming as the speaker for the current frame of image data.

On the other hand, if it is determined at step S144 that there is more than one participant having a head at a position which corresponds to the direction from which the speech is coming, then, at step S148, archive processor 58 determines whether one of those participants was identified as the speaker in the preceding frame of

image data.

5 If it is determined at step S148 that one of the participants in the direction from which the speech is coming was selected as the speaker in the preceding frame of image data, then, at step S150, archive processor 58 selects the speaker identified for the previous frame of image data as the speaker for the current frame of image data, too. This is because is it likely that the speaker
10 in the previous frame of image data is the same as the speaker in the current frame.

On the other hand, if it is determined at step S148 that none of the participants in the direction from which the speech is coming is the participant identified as the speaker in the preceding frame, or if no speaker was
15 identified for the previous frame, then, at step S152, archive processor 58 selects each of the participants in the direction from which the speech is coming as a "potential" speaking participant.
20

Referring again to Figure 8, at step S86, archive processor 58 stores the viewing parameter value for each speaking participant, that is the viewing parameter value
25 generated at step S80 defining at whom or what each

speaking participant is looking, for subsequent analysis, for example in memory 42.

At step S88, archive processor 58 informs voice
5 recognition processor 54 of the identity of each speaking
participant determined at step S84. In response, voice
recognition processor 54 selects the speech recognition
parameters for the speaking participant(s) from speech
recognition parameter store 56 and uses the selected
10 parameters to perform speech recognition processing on
the received speech data to generate text data
corresponding to the words spoken by the speaking
participant(s).

15 On the other hand, if it is determined at step S82 that
the received sound data does not contain any speech, then
steps S84 to S88 are omitted.

At step S89, archive processor 58 determines which image
20 data is to be stored in the meeting archive database 60,
that is, the image data from which of the cameras 2-1, 2-
2 and 2-3 is to be stored.

Figure 13 shows the processing operations performed by
25 archive processor 58 at step S89.

Referring to Figure 13, at step S160, archive processor 58 determines whether any speech was detected at step S82 (Figure 8) for the current frame of image data.

5 If it is determined at step S160 that there is no speech for the current frame, then, at step S162, archive processor 58 selects a default camera as the camera from which image data is to be stored. More particularly, in this embodiment, archive processor 58 selects the camera or cameras from which image data was recorded for the previous frame, or, if the current frame being processed is the very first frame, then archive processor 58 selects one of the cameras 2-1, 2-2, 2-3 at random.

10 On the other hand, if it is determined at step S160 that there is speech for the current frame being processed then, at step S164, archive processor 58 reads the viewing parameter previously stored at step S86 for the next speaking participant (this being the first speaking participant the first time step S164 is performed) to determine the person or object at which that speaking participant is looking.

25 At step S166, archive processor 58 reads the head position and orientation (determined at step S44 in

Figure 5) for the speaking participant currently being considered, together with the head position and orientation of the participant at which the speaking participant is looking (determined at step S44 in Figure 5) or the position and orientation of the object at which the speaking participant is looking (stored at step S26 in Figure 3).

At step S168 archive processor 58 processes the positions and orientations read at step S166 to determine which of the cameras 2-1, 2-2, 2-3 best shows both the speaking participant and the participant or object at which the speaking participant is looking, and selects this camera as a camera from which image data for the current frame is to be stored in meeting archive database 60.

Figure 14 shows the processing operations performed by archive processor 58 at step S168.

Referring to Figure 14, at step S176, archive processor 58 reads the three-dimensional position and viewing direction of the next camera (this being the first camera the first time step S176 is performed), this information having previously been generated and stored at step S24 in Figure 3.

At step S178, archive processor 58 uses the information read at step S176 together with information defining the three-dimensional head position and orientation of the speaking participant (determined at step S44 in Figure 5) and the three-dimensional head position and orientation of the participant at whom the speaking participant is looking (determined at step S44 in Figure 5) or the three-dimensional position and orientation of the object being looked at (stored at step S26 in Figure 3) to determine whether the speaking participant and the participant or object at which the speaking participant is looking are both within the field of view of the camera currently being considered (that is, whether the camera currently being considered can see both the speaking participant and the participant or object at which the speaking participant is looking). More particularly, in this embodiment, archive processor 58 evaluates the following equations and determines that the camera can see both the speaking participant and the participant or object at which the speaking participant is looking if all of the inequalities hold:

$$\left| \arccos \left[\frac{1}{\sqrt{(X_{pl} - X_c)^2 + (Y_{pl} - Y_c)^2}} \begin{pmatrix} X_{pl} - X_c \\ Y_{pl} - Y_c \end{pmatrix} \cdot \begin{pmatrix} dX_c \\ dY_c \end{pmatrix} \right] \right| < \theta_h$$

....(1)

$$\left| \arccos \left[\frac{1}{\sqrt{(X_{p1}-X_c)^2+(Y_{p1}-Y_c)^2+(Z_{p1}-Z_c)^2}} \begin{pmatrix} X_{p1}-X_c \\ Y_{p1}-Y_c \\ Z_{p1}-Z_c \end{pmatrix} \cdot \begin{pmatrix} dX_c \\ dY_c \\ dZ_c \end{pmatrix} \right] \right| < \theta_v$$

.....(2)

5

$$\left| \arccos \left[\frac{1}{\sqrt{(X_{p2}-X_c)^2+(Y_{p2}-Y_c)^2+(Z_{p2}-Z_c)^2}} \begin{pmatrix} X_{p2}-X_c \\ Y_{p2}-Y_c \\ Z_{p2}-Z_c \end{pmatrix} \cdot \begin{pmatrix} dX_c \\ dY_c \\ dZ_c \end{pmatrix} \right] \right| < \theta_h$$

.....(3)

10

$$\left| \arccos \left[\frac{1}{\sqrt{(X_{p2}-X_c)^2+(Y_{p2}-Y_c)^2+(Z_{p2}-Z_c)^2}} \begin{pmatrix} X_{p2}-X_c \\ Y_{p2}-Y_c \\ Z_{p2}-Z_c \end{pmatrix} \cdot \begin{pmatrix} dX_c \\ dY_c \\ dZ_c \end{pmatrix} \right] \right| < \theta_v$$

.....(4)

15

where:

(X_c , Y_c , Z_c) are the x, y and z coordinates respectively of the principal point of the camera (previously determined and stored at step S24 in Figure 3)

20

(dX_c , dY_c , dZ_c) represent the viewing direction of the camera in the x, y and z directions respectively (again determined and stored at step S24 in Figure 3)

25

002220" 66926560

θ_h and θ_v are the angular fields of view of the camera in the horizontal and vertical directions respectively (again determined and stored at step S24 in Figure 3)

5

(X_{p1}, Y_{p1}, Z_{p1}) are the x, y and z coordinates respectively of the centre of the head of the speaking participant (determined at step S44 in Figure 5)

10

$(dX_{p1}, dY_{p1}, dZ_{p1})$ represent the orientation of the viewing ray 90 of the speaking participant (again determined at step S44 in Figure 5)

15

(X_{p2}, Y_{p2}, Z_{p2}) are the x, y and z coordinates respectively of the centre of the head of the person at whom the speaking participant is looking (determined at step S44 in Figure 5) or of the centre of the surface of the object at which the speaking participant is looking (determined at step S26 in Figure 3)

20

$(dX_{p2}, dY_{p2}, dZ_{p2})$ represent the direction in the x, y and z directions respectively of the viewing ray 90 of the participant at whom the speaking

25

00220" E52E560

participant is looking (again determined at step S44 in Figure 5) or of the normal to the object surface at which the speaking participant is looking (determined at step S26 in Figure 3).

5

If it is determined at step S178 that the camera can see both the speaking participant and the participant or object at which the speaking participant is looking (that is, the inequalities in each of equations (1), (2), (3) and (4) above hold), then, at step S180, archive processor 58 calculates and stores a value representing the quality of the view that the camera currently being considered has of the speaking participant. More particularly, in this embodiment, archive processor 58 calculates a quality value, Q1, using the following equation:

10

15

$$Q1 = \frac{1}{\sqrt{(X_c - X_{pl})^2 + (Y_c - Y_{pl})^2 + (Z_c - Z_{pl})^2}} \begin{pmatrix} X_c - X_{pl} \\ Y_c - Y_{pl} \\ Z_c - Z_{pl} \end{pmatrix} \cdot \begin{pmatrix} dX_{pl} \\ dY_{pl} \\ dZ_{pl} \end{pmatrix} \dots (5)$$

20

where the definitions of the terms are the same as those given for equations (1) and (2) above.

The quality value, Q1, calculated at step S180 is a

scalar, having a value between -1 and +1, with the value being -1 if the back of the speaking participant's head is directly facing the camera, +1 if the face of the speaking participant is directly facing the camera, and a value in-between for other orientations of the speaking participant's head.

At step S182, archive processor 58 calculates and stores a value representing the quality of the view that the camera currently being considered has of the participant or object at which the speaking participant is looking. More particularly, in this embodiment, archive processor 58 calculates a quality value, Q2, using the following equation:

$$Q2 = \frac{1}{\sqrt{(X_c - X_{p2})^2 + (Y_c - Y_{p2})^2 + (Z_c - Z_{p2})^2}} \begin{pmatrix} X_c - X_{p2} \\ Y_c - Y_{p2} \\ Z_c - Z_{p2} \end{pmatrix} \cdot \begin{pmatrix} dX_{p2} \\ dY_{p2} \\ dZ_{p2} \end{pmatrix} \dots (6)$$

where the definitions of the parameters are the same as those given for equations (3) and (4) above.

Again, Q2 is a scalar having a value between -1 if the back of the head of the participant or the back of the surface of the object is directly facing the camera, +1

5

15

20

25

received. Steps S176 to S186 are repeated until the processing described above has been performed for each camera.

5 At step S188, archive processor 58 compares the "worst view" values stored for each of the cameras when processing was performed at step S184 (that is, the value of Q1 or Q2 stored for each camera at step S184) and selects the highest one of these stored values. This
10 highest value represents the "best worst view" and accordingly, at step S188, archive processor 58 selects the camera for which this "best worst view" value was stored at step S184 as a camera from which image data should be stored in the meeting archive database, since
15 this camera has the best view of both the speaking participant and the participant or object at which the speaking participant is looking.

At step S170, archive processor 58 determines whether
20 there is another speaking participant, including any "potential" speaking participants. Steps S164 to S170 are repeated until the processing described above has been performed for each speaking participant and each "potential" speaking participant.

Referring again to Figure 8, at step S90, archive processor 58 encodes the current frame of video data received from the camera or cameras selected at step S89 and the sound data received from microphone array 4 as MPEG 2 data in a conventional manner, and stores the encoded data in meeting archive database 60.

Figure 15 schematically illustrates the storage of data in meeting archive database 60. The storage structure shown in Figure 15 is notional and is provided to assist understanding by illustrating the links between the stored information; it does not necessarily represent the exact way in which data is stored in the memory comprising meeting archive database 60.

Referring to Figure 15, meeting archive database 60 stores time information represented by the horizontal axis 200, on which each unit represents a predetermined amount of time, for example the time period of one frame of video data received from a camera. (It will, of course, be appreciated that the meeting archive database 60 will generally contain many more time units than the number shown in Figure 15.) The MPEG 2 data generated at step S90 is stored as data 202 in meeting archive database 60, together with timing information (this

timing information being schematically represented in Figure 15 by the position of the MPEG 2 data 202 along the horizontal axis 200).

5 Referring again to Figure 8, at step S92, archive processor 58 stores any text data generated by voice recognition processor 54 at step S88 for the current frame in meeting archive database 60 (indicated at 204 in Figure 15). More particularly, the text data is stored with a link to the corresponding MPEG 2 data, this link being represented in Figure 15 by the text data being stored in the same vertical column as the MPEG 2 data. As will be appreciated, there will not be any text data for storage from participants who are not speaking. In the example shown in Figure 15, text is stored for the first ten time slots for participant 1 (indicated at 206), for the twelfth to twentieth time slots for participant 3 (indicated at 208), and for the twenty-first time slot for participant 4 (indicated at 210). No text is stored for participant 2 since, in this example, participant 2 did not speak during the time slots shown in Figure 15.

At step S94, archive processor 58 stores the viewing parameter value generated for the current frame for each

participant at step S80 in the meeting archive database
60 (indicated at 212 in Figure 15). Referring to Figure
15, a viewing parameter value is stored for each
participant together with a link to the associated MPEG
2 data 202 and the associated text data 204 (this link
5 being represented in Figure 15 by the viewing parameter
values being shown in the same column as the associated
MPEG 2 data 202 and associated text data 204). Thus,
referring to the first time slot in Figure 15 by way of
10 example, the viewing parameter value for participant 1 is
3, indicating that participant 1 is looking at
participant 3, the viewing parameter value for
participant 2 is 5, indicating that participant 2 is
looking at the flip chart 14, the viewing parameter value
15 for participant 3 is 1, indicating that participant 3 is
looking at participant 1, and the viewing parameter value
for participant 4 is "0", indicating that participant 4
is not looking at any of the other participants (in the
example shown in Figure 1, the participant indicated at
20 12 is looking at her notes rather than any of the other
participants).

At step S96, central controller 36 and archive processor
58 determine whether one of the participants in the
25 meeting has stopped speaking. In this embodiment, this

check is performed by examining the text data 204 to determine whether text data for a given participant was present for the previous time slot, but is not present for the current time slot. If this condition is satisfied for any participant (that is, a participant has stopped speaking), then, at step S98, archive processor 58 processes the viewing parameter values previously stored when step S86 was performed for each participant who has stopped speaking (these viewing parameter values defining at whom or what the participant was looking during the period of speech which has now stopped) to generate data defining a viewing histogram. More particularly, the viewing parameter values for the period in which the participant was speaking are processed to generate data defining the percentage of time during that period that the speaking participant was looking at each of the other participants and objects.

Figures 16A and 16B show the viewing histograms corresponding to the periods of text 206 and 208 respectively in Figure 15.

Referring to Figure 15 and Figure 16A, during the period 206 when participant 1 was speaking, he was looking at participant 3 for six of the ten time slots (that is, 60%

of the total length of the period for which he was talking), which is indicated at 300 in Figure 16A, and at participant 4 for four of the ten time slots (that is, 40% of the time), which is indicated at 310 in Figure 16A.

Similarly, referring to Figure 15 and Figure 16B, during the period 208, participant 3 was looking at participant 1 for approximately 45% of the time, which is indicated at 320 in Figure 16B, at object 5 (that is, the flip chart 14) for approximately 33% of the time, indicated at 330 in Figure 16B, and at participant 2 for approximately 22% of the time, which is indicated at 340 in Figure 16B.

Referring again to Figure 8, at step S100, each viewing histogram generated at step S98 is stored in the meeting archive database 60 linked to the associated period of text for which it was generated. Referring to Figure 15, the stored viewing histograms are indicated at 214, with the data defining the histogram for the text period 206 indicated at 216, and the data defining the histogram for the text period 208 indicated at 218. In Figure 15, the link between the viewing histogram and the associated text is represented by the viewing histogram being stored in the same columns as the text data.

On the other hand, if it is determined at step S96 that, for the current time period, one of the participants has not stopped speaking, then steps S98 and S100 are omitted.

5

At step S102, archive processor 58 corrects data stored in the meeting archive database 60 for the previous frame of video data (that is, the frame preceding the frame for which data has just been generated and stored at steps S80 to S100) and other preceding frames, if such correction is necessary.

10

Figure 17 shows the processing operations performed by archive processor 58 at step S102.

15

Referring to Figure 17, at step S190, archive processor 58 determines whether any data for a "potential" speaking participant is stored in the meeting archive database 60 for the next preceding frame (this being the frame which immediately precedes the current frame the first time step S190 is performed, that is the "i-1"th frame if the current frame is the "i"th frame).

20

If it is determined at step S190 that no data is stored for a "potential" speaking participant for the preceding

25

002220" E526560

frame being considered, then it is not necessary to correct any data in the meeting archive database 60.

On the other hand, if it is determined at step S190 that data for a "potential" speaking participant is stored for the preceding frame being considered, then, at step S192, archive processor 58 determines whether one of the "potential" speaking participants for which data was stored for the preceding frame is the same as a speaking participant (but not a "potential" speaking participant) identified for the current frame, that is a speaking participant identified at step S146 in Figure 12.

If it is determined at step S192 that none of the "potential" speaking participants for the preceding frame is the same as a speaking participant identified at step S146 for the current frame, then no correction of the data stored in the meeting archive database 60 for the preceding frame being considered is carried out.

On the other hand, if it is determined at step S192 that a "potential" speaking participant for the preceding frame is the same as a speaking participant identified at step S146 for the current frame, then, at step S194, archive processor 58 deletes the text data 204 for the

[illegible]

10

15

20

25

at step S192 that none of the "potential" speaking participants in the preceding frame being considered is the same as a speaking participant unambiguously identified for the current frame. In this way, in cases where "potential" speaking participants were identified for a number of successive frames, the data stored in the meeting archive database is corrected if the actual speaking participant from among the "potential" speaking participants is identified in the next frame.

Referring again to Figure 8, at step S104, central controller 36 determines whether another frame of video data has been received from the cameras 2-1, 2-2, 2-3. Steps S80 to S104 are repeatedly performed while image data is received from the cameras 2-1, 2-2, 2-3.

When data is stored in meeting archive database 60, then the meeting archive database 60 may be interrogated to retrieve data relating to the meeting.

Figure 18 shows the processing operations performed to search the meeting archive database 60 to retrieve data relating to each part of the meeting which satisfies search criteria specified by a user.

Referring to Figure 18, at step S200, central controller 36 causes display processor 64 to display a message on display device 26 requesting the user to enter information defining the search of meeting archive database 60 which is required. More particularly, in this embodiment, central controller 100 causes the display shown in Figure 19A to appear on display device 26.

Referring to Figure 19A, the user is requested to enter information defining the part or parts of the meeting which he wishes to find in the meeting archive database 60. More particularly, in this embodiment, the user is requested to enter information 400 defining a participant who was talking, information 410 comprising one or more key words which were said by the participant identified in information 400, and information 420 defining the participant or object at which the participant identified in information 400 was looking when he was talking. In addition, the user is able to enter time information defining a portion or portions of the meeting for which the search is to be carried out. More particularly, the user can enter information 430 defining a time in the meeting beyond which the search should be discontinued (that is, the period of the meeting before the specified

time should be searched), information 440 defining a time in the meeting after which the search should be carried out, and information 450 and 460 defining a start time and end time respectively between which the search is to be carried out. In this embodiment, information 430, 440, 450 and 460 may be entered either by specifying a time in absolute terms, for example in minutes, or in relative terms by entering a decimal value which indicates a proportion of the total meeting time. For example, entering the value 0.25 as information 430 would restrict the search to the first quarter of the meeting.

In this embodiment, the user is not required to enter all of the information 400, 410 and 420 for one search, and instead may omit one or two pieces of this information. If the user enters all of the information 400, 410 and 420, then the search will be carried out to identify each part of the meeting in which the participant identified in information 400 was talking to the participant or object identified in information 420 and spoke the key words defined in information 410. On the other hand, if information 410 is omitted, then a search will be carried out to identify each part of the meeting in which the participant defined in information 400 was talking to the participant or object defined in information 420

irrespective of what was said. If information 410 and 420 is omitted, then a search is carried out to identify each part of the meeting in which the participant defined in information 400 was talking, irrespective of what was said and to whom. If information 400 is omitted, then a search is carried out to identify each part of the meeting in which any of the participants spoke the key words defined in information 410 while looking at the participant or object defined in information 420. If information 400 and 410 is omitted, then a search is carried out to identify each part of the meeting in which any of the participants spoke to the participant or object defined in information 420. If information 420 is omitted, then a search is carried out to identify each part of the meeting in which the participant defined in information 400 spoke the key words defined in information 410, irrespective of to whom the key words were spoken. Similarly, if information 400 and 420 is omitted, then a search is carried out to identify each part of the meeting in which the key words identified in information 410 were spoken, irrespective of who said the key words and to whom.

In addition, the user may enter all of the time information 430, 440, 450 and 460 or may omit one or more

pieces of this information.

Further, known Boolean operators and search algorithms may be used in combination with key words entered in information 410 to enable the searcher to search for combinations or alternatives of words.

Once the user has entered all of the required information to define the search, he begins the search by clicking on area 470 using a user input device such as the mouse 30.

Referring again to Figure 18, at step S202, the search information entered by the user is read by central controller 36 and the instructed search is carried out. More particularly, in this embodiment, central controller 36 converts any participant or object names entered in information 400 or 420 to identification numbers using the table 80 (Figure 4), and considers the text information 204 for the participant defined in information 400 (or all participants if information 400 is not entered). If information 420 has been entered by the user, then, for each period of text, central controller 36 checks the data defining the corresponding viewing histogram to determine whether the percentage of viewing time in the histogram for the participant or

object defined in information 420 is equal to or above a threshold, which, in this embodiment, is 25%. In this way, periods of speech (text) are considered to satisfy the criteria that a participant defined in information 400 was talking to the participant or object defined in information 420 even if the speaking participant looked at other participants or objects while speaking, provided that the speaking participant looked at the participant or object defined in information 420 for at least 25% of the time of the speech. Thus, for example, a period of speech in which the value of the viewing histogram is equal to or above 25% for two or more participants would be identified if any of these participants were specified in information 420. If the information 410 has been input by the user, then central controller 36 and text searcher 62 search each portion of text previously identified on the basis of information 400 and 420 (or all portions of text if information 400 and 420 was not entered) to identify each portion containing the key word(s) identified in information 410. If any time information has been entered by the user, then the searches described above are restricted to the meeting times defined by those limits.

At step S204, central controller 36 causes display

processor 64 to display a list of relevant speeches identified during the search to the user on display device 26. More particularly, central controller 36 causes information such as that shown in Figure 19B to be displayed to the user. Referring to Figure 19B, a list is produced of each speech which satisfies the search parameters, and information is displayed defining the start time for the speech both in absolute terms and as a proportion of the full meeting time. The user is then able to select one of the speeches for playback, for example by clicking on the required speech in the list using the mouse 30.

At step S206, central controller 36 reads the selection made by the user at step S204, and plays back the stored MPEG 2 data 202 for the relevant part of the meeting from meeting archive database 60. More particularly, central controller 36 and display processor 64 decode the MPEG 2 data 202 and output the image data and sound via display device 26. If image data from more than one camera is stored for part, or the whole, of the speech to be played back, then this is indicated to the user on display device 26 and the user is able to select the image data which is to be replayed by inputting instructions to central controller 36, for example using keyboard 28.

At step S208, central controller 36 determines whether the user wishes to cease interrogating the meeting archive database 60 and, if not, steps S200 to S208 are repeated.

5

Various modifications and changes can be made to the embodiment of the invention described above.

In the embodiment above, at step S34 (Figure 3) and step S70 (Figure 7) the head of each of the participants in the meeting is tracked. In addition, however, objects for which data was stored at step S4 and S26 could also be tracked if they moved (such objects may comprise, for example, notes which are likely to be moved by a participant or an object which is to be passed between the participants).

In the embodiment above, image data is processed from a plurality of video cameras 2-1, 2-2, 2-3. However, instead, image data may be processed from a single video camera. In this case, for example, only step S42-1 (Figure 5) is performed and steps S42-2 to S42-n are omitted. Similarly, step S44 is omitted and the 3D position and orientation of the participant's head for the current frame of image data are taken to be the 3D

position and orientation determined at step S58 (Figure 6) during the processing performed at step S42-1. At step S46, the position for the head features input to the Kalman filter would be the position in the image data from the single camera. Further, step S89 (Figure 8) to select the camera from which image data is to be recorded in the meeting archive database 60 would also be omitted.

In the embodiment above, at step S168 (Figure 13), processing is performed to identify the camera which has the best view of the speaking participant and also the participant or object at which the speaking participant is looking. However, instead of identifying the camera in the way described in the embodiment above, it is possible for a user to define during the initialisation of processing apparatus 24 which of the cameras 2-1, 2-2, 2-3 has the best view of each respective pair of the seating positions around the meeting table and/or the best view of each respective seating position and a given object (such as flip chart 14). In this way, if it is determined that the speaking participant and the participant at whom the speaking participant is looking are in predefined seating positions, then the camera defined by the user to have the best view of those predefined seating positions can be selected as a camera

from which image data is to be stored. Similarly, if the speaking participant is in a predefined position and is looking at an object, then the camera defined by the user to have the best view of that predefined seating position and object can be selected as the camera from which image data is to be stored.

In the embodiment above, at step S162 (Figure 13) a default camera is selected as a camera from which image data was stored for the previous frame. Instead, however, the default camera may be selected by a user, for example during the initialisation of processing apparatus 24.

In the embodiment above, at step S194 (Figure 17), the text data 204 is deleted from meeting archive database 60 for the "potential" speaking participants who have now been identified as actually not being speaking participants. In addition, however, the associated viewing histogram data 214 may also be deleted. Further, if MPEG 2 data 202 from more than one of the cameras 2-1, 2-2, 2-3 was stored, then the MPEG 2 data related to the "potential" speaking participants may also be deleted.

In the embodiment above, when it is not possible to

uniquely identify a speaking participant, "potential" speaking participants are defined, data is processed and stored in meeting archive database 60 for the potential speaking participants, and subsequently the data stored in the meeting archive database 60 is corrected (step S102 in Figure 8). However, instead, rather than processing and storing data for "potential" speaking participants, video data received from cameras 2-1, 2-2 and 2-3 and audio data received from microphone array 4 may be stored for subsequent processing and archiving when the speaking participant has been identified from data relating to a future frame. Alternatively, when the processing performed at step S144 (Figure 12) results in an indication that there is more than one participant in the direction from which the speech is coming, image data from the cameras 2-1, 2-2 and 2-3 may be processed to detect lip movements of the participants and to select as the speaking participant the participant in the direction from which the speech is coming whose lips are moving.

In the embodiment above, processing is performed to determine the position of each person's head, the orientation of each person's head and a viewing parameter for each person defining at whom or what the person is looking. The viewing parameter value for each person is

then stored in the meeting archive database 60 for each frame of image data. However, it is not necessary to determine a viewing parameter for all of the people. For example, it is possible to determine a viewing parameter for just the speaking participant, and to store just this viewing parameter value in the meeting archive database 60 for each frame of image data. Accordingly, in this case, it would be necessary to determine the orientation of only the speaking participant's head. In this way, processing requirements and storage requirements can be reduced.

In the embodiment above, at step S202 (Figure 18), the viewing histogram for a particular portion of text is considered and it is determined that the participant was talking to a further participant or object if the percentage of gaze time for the further participant or object in the viewing histogram is equal to or above a predetermined threshold. Instead, however, rather than using a threshold, the participant or object at whom the speaking participant was looking during the period of text (speech) may be defined to be the participant or object having the highest percentage gaze value in the viewing histogram (for example participant 3 in Figure 16A, and participant 1 in Figure 16B).

In the embodiment above, the MPEG 2 data 202, the text data 204, the viewing parameters 212 and the viewing histograms 214 are stored in meeting archive database 60 in real time as data is received from cameras 2-1, 2-2 and 2-3 and microphone array 4. However, instead, the video and sound data may be stored and data 202, 204, 212 and 214 generated and stored in meeting archive database 60 in non-real-time.

In the embodiment above, the MPEG 2 data 202, the text data 204, the viewing parameters 212 and the viewing histograms 214 are generated and stored in the meeting archive database 60 before the database is interrogated to retrieve data for a defined part of the meeting. However, some, or all, of the viewing histogram data 214 may be generated in response to a search of the meeting archive database 60 being requested by the user by processing the data already stored in meeting archive database 60, rather than being generated and stored prior to such a request. For example, although in the embodiment above the viewing histograms 214 are calculated and stored in real-time at steps S98 and S100 (Figure 8), these histograms could be calculated in response to a search request being input by the user.

In the embodiment above, text data 204 is stored in meeting archive database 60. Instead, audio data may be stored in the meeting archive database 60 instead of the text data 204. The stored audio data would then either
5 itself be searched for key words using voice recognition processing or converted to text using voice recognition processing and the text search using a conventional text searcher.

10 In the embodiment above, processing apparatus 24 includes functional components for receiving and generating data to be archived (for example, central controller 36, head tracker 50, head model store 52, direction processor 53, voice recognition processor 54, speech recognition
15 parameter store 56 and archive processor 58), functional components for storing the archive data (for example meeting archive database 60), and also functional components for searching the database and retrieving information therefrom (for example central controller 36
20 and text searcher 62). However, these functional components may be provided in separate apparatus. For example, one or more apparatus for generating data to be archived, and one or more apparatus for database searching may be connected to one or more databases via
25 a network, such as the Internet.

Also, referring to Figure 20, video and sound data from one or more meetings 500, 510, 520 may be input to a data processing and database storage apparatus 530 (which comprises functional components to generate and store the archive data), and one or more database interrogation apparatus 540, 550 may be connected to the data processing and database storage apparatus 530 for interrogating the database to retrieve information therefrom.

In the embodiment above, processing is performed by a computer using processing routines defined by programming instructions. However, some, or all, of the processing could be performed using hardware.

Although the embodiment above is described with respect to a meeting taking place between a number of participants, the invention is not limited to this application, and, instead, can be used for other applications, such as to process image and sound data on a film set etc.

Different combinations of the above modifications are, of course, possible and other changes and modifications can be made without departing from the spirit and scope of

the invention.

Second Embodiment

5 Referring to Figure 21, a video camera 602 and one or more microphones 604 are used to record image data and sound data respectively from a meeting taking place between a group of people 606, 608, 610, 612.

10 The image data from the video camera 602 and the sound data from the microphones 604 is input via cables (not shown) to a computer 620 which processes the received data and stores data in a database to create an archive record of the meeting from which information can be
15 subsequently retrieved.

Computer 620 comprises a conventional personal computer having a processing apparatus 624 containing, in a conventional manner, one or more processors, memory,
20 sound card etc., together with a display device 626 and user input devices, which, in this embodiment, comprise a keyboard 628 and a mouse 630.

The components of computer 620 and the input and output
25 of data therefrom are schematically shown in Figure 22.

Referring to Figure 22, the processing apparatus 624 is programmed to operate in accordance with programming instructions input, for example, as data stored on a data storage medium, such as disk 632, and/or as a signal 634 input to the processing apparatus 624, for example from a remote database, by transmission over a communication network (not shown) such as the Internet or by transmission through the atmosphere, and/or entered by a user via a user input device such as keyboard 628 or other input device.

When programmed by the programming instructions, processing apparatus 624 effectively becomes configured into a number of functional units for performing processing operations. Examples of such functional units and their interconnections are shown in Figure 22. The illustrated units and interconnections in Figure 22 are, however, notional and are shown for illustration purposes only, to assist understanding; they do not necessarily represent the exact units and connections into which the processor, memory etc of the processing apparatus 624 become configured.

Referring to the functional units shown in Figure 22, a central controller 636 processes inputs from the user

input devices 628, 630 and receives data input to the processing apparatus 624 by a user as data stored on a storage device, such as disk 638, or as a signal 640 transmitted to the processing apparatus 624. The central controller 636 also provides controller and processing for a number of the other functional units. Memory 642 is provided for use by central controller 636 and other functional units.

Head tracker 650 processes the image data received from video camera 602 to track the position and orientation in three dimensions of the head of each of the participants 606, 608, 610, 612 in the meeting. In this embodiment, to perform this tracking, head tracker 650 uses data defining a three-dimensional computer model of the head of each of the participants and data defining features thereof which is stored in head model store 652, as will be described below.

Voice recognition processor 654 processes sound data received from microphones 604. Voice recognition processor 640 operates in accordance with a conventional voice recognition program, such as "Dragon Dictate" or IBM "ViaVoice", to generate text data corresponding to the words spoken by the participants 606, 608, 610, 612.

To perform the voice recognition processing, voice recognition processor 654 uses data defining the speech recognition parameters for each participant 606, 608, 610, 612, which is stored in speech recognition parameter store 656. More particularly, the data stored in speech recognition parameter store 656 comprises data defining the voice profile of each participant which is generated by training the voice recognition processor in a conventional manner. For example, the data comprises the data stored in the "user files" of Dragon Dictate after training.

Archive processor 658 generates data for storage in meeting archive database 660 using data received from head tracker 650 and voice recognition processor 654. More particularly, as will be described below, the video data from camera 602 and sound data from microphones 604 is stored in meeting archive database 660 together with text data from voice recognition processor 654 and data defining at whom each participant in the meeting was looking at a given time.

Text searcher 662, in conjunction with central controller 636, is used to search the meeting archive database 660 to find and replay the sound and video data for one or

more parts of the meeting which meet search criteria specified by a user, as will be described in further detail below.

5 Display processor 664 under control of central controller 636 displays information to a user via display device 626 and also replays sound and video data stored in meeting archive database 660.

10 Output processor 666 outputs part or all of the data from archive database 660, for example on a storage device such as disk 668 or as a signal 670.

15 Before beginning the meeting, it is necessary to initialise computer 620 by entering data which is necessary to enable processing apparatus 624 to perform the required processing operations.

20 Figure 23 shows the processing operations performed by processing apparatus 624 during this initialisation.

25 Referring to Figure 23, at step S302, central controller 636 causes display processor 664 to display a message on display device 626 requesting the user to input the names of each person who will participate in

the meeting.

At step S304, upon receipt of data defining the names, for example input by the user using keyboard 628, central controller 636 allocates a unique participant number to each participant, and stores data, for example table 680 shown in Figure 24, defining the relationship between the participant numbers and the participants' names in the meeting archive database 660.

At step S306, central controller 636 searches the head model store 652 to determine whether data defining a head model is already stored for each participant in the meeting.

If it is determined at step S306 that a head model is not already stored for one or more of the participants, then, at step S308, central controller 636 causes display processor 664 to display a message on display device 626 requesting the user to input data defining a head model of each participant for whom a model is not already stored.

In response, the user enters data, for example on a storage medium such as disk 638 or by downloading the

data as a signal 640 from a connected processing apparatus, defining the required head models. Such head models may be generated in a conventional manner, for example as described in "An Analysis/Synthesis Cooperation for Head Tracking and Video Face Cloning" by Valente et al in Proceedings ECCV '98 Workshop on Perception of Human Action, University of Freiberg, Germany, June 6 1998.

At step S310, central controller 636 stores the data input by the user in head model store 652.

At step S312, central controller 636 and display processor 664 render each three-dimensional computer head model input by the user to display the model to the user on display device 626, together with a message requesting the user to identify at least seven features in each model.

In response, the user designates using mouse 630 points in each model which correspond to prominent features on the front, sides and, if possible, the back, of the participant's head, such as the corners of eyes, nostrils, mouth, ears or features on glasses worn by the participant, etc.

At step S314, data defining the features identified by the user is stored by central controller 636 in head model store 652.

5 On the other hand, if it is determined at step S306 that a head model is already stored in head model store 652 for each participant, then steps S308 to S314 are omitted.

10 At step S316, central controller 636 searches speech recognition parameter store 656 to determine whether speech recognition parameters are already stored for each participant.

15 If it is determined at step S316 that speech recognition parameters are not available for all of the participants, then, at step S318, central controller 636 causes display processor 664 to display a message on display device 626 requesting the user to input the speech recognition
20 parameters for each participant for whom the parameters are not already stored.

In response, the user enters data, for example on a storage medium such as disk 638 or as a signal 640 from
25 a remote processing apparatus, defining the necessary

speech recognition parameters. As noted above, these parameters define a profile of the user's speech and are generated by training a voice recognition processor in a conventional manner. Thus for example, in the case of a voice recognition processor comprising Dragon Dictate, the speech recognition parameters input by the user correspond to the parameters stored in the "user files" of Dragon Dictate.

At step S320, the data input by the user is stored by central controller 636 in the speech recognition parameter store 656.

On the other hand, if it is determined at step S316 that the speech recognition parameters are already available for each of the participants, then steps S318 and S320 are omitted.

At step S322, central controller 636 causes display processor 664 to display a message on display device 626 requesting the user to perform steps to enable the camera 602 to be calibrated.

In response, the user carries out the necessary steps and, at step S324, central controller 636 performs

processing to calibrate the camera 602. More particularly, in this embodiment, the steps performed by the user and the processing performed by central controller 636 are carried out in a manner such as that described in "Calibrating and 3D Modelling with a Multi-Camera System" by Wiles and Davison in 1999 IEEE Workshop on Multi-View Modelling and Analysis of Visual Scenes, ISBN 0769501109. This generates calibration data defining the position and orientation of the camera 602 with respect to the meeting room and also the intrinsic camera parameters (aspect ratio, focal length, principal point, and first order radial distortion coefficient). The calibration data is stored in memory 642.

At step S326, central controller 636 causes display processor 664 to display a message on display device 626 requesting the next participant in the meeting (this being the first participant the first time step S326 is performed) to sit down.

At step S328, processing apparatus 624 waits for a predetermined period of time to give the requested participant time to sit down, and then, at step S330, central controller 636 processes image data from camera 602 to determine an estimate of the position of the

00220" 032200

5 seated participant's head. More particularly, in this embodiment, central controller 636 carries out processing in a conventional manner to identify each portion in a frame of image data from camera 622 which has a colour corresponding to the colour of the skin of the participant (this colour being determined from the data defining the head model of the participant stored in head model store 652), and then selects the portion which corresponds to the highest position in the meeting room (since it is assumed that the head will be the highest skin-coloured part of the body). Using the position of the identified portion in the image and the camera calibration parameters determined at step S324, central controller 636 then determines an estimate of the three-dimensional position of the head in a conventional manner.

20 At step S332, central controller 636 determines an estimate of the orientation of the participant's head in three dimensions. More particularly, in this embodiment, central controller 636 renders the three-dimensional computer model of the participant's head stored in head model store 652 for a plurality of different orientations of the model to produce a respective two-dimensional image of the model for each orientation, compares each

25

two-dimensional image of the model with the part of the videoframe from camera 602 which shows the participant's head, and selects the orientation for which the image of the model best matches the video image data. In this embodiment, the computer model of the participant's head is rendered in 108 different orientations to produce image data for comparing with the video data from camera 602. These orientations correspond to 36 rotations of the head model in 10° steps for each of three head inclinations corresponding to 0° (looking straight ahead), +45° (looking up) and -45° (looking down). When comparing the image data produced by rendering the head model with the video data from camera 602, a conventional technique is used, for example as described in "Head Tracking Using a Textured Polygonal Model" by Schödl, Haro & Essa in Proceedings 1998 Workshop on Perceptual User Interfaces.

At step S334, the estimate of the position of the participant's head generated at step S330 and the estimate of the orientation of the participant's head generated at step S332 are input to head tracker 650 and frames of image data received from camera 602 are processed to track the head of the participant. More particularly, in this embodiment, head tracker 650

performs processing to track the head in a conventional manner, for example as described in "An Analysis/Synthesis Cooperation for Head Tracking and Video Face Cloning" by Valente et al in Proceedings EECV '98 Workshop on Perception of Human Action, University of Freiberg, Germany, June 6 1998.

Figure 25 summarises the processing operations performed by head tracker 650 at step S334.

Referring to Figure 25, at step S350, head tracker 650 reads the current estimates of the 3D position and orientation of the participant's head, these being the estimates produced at steps S330 and S332 in Figure 23 the first time step S350 is performed.

At step S352, head tracker 650 uses the camera calibration data generated at step S324 to render the three-dimensional computer model of the participant's head stored in head model store 652 in accordance with the estimates of position and orientation read at step S350.

At step S354, head tracker 650 processes the image data for the current frame of video data received from

camera 602 to extract the image data from each area which surrounds the expected position of one of the head features identified by the user and stored at step S314, the expected positions being determined from the estimates read at step S350 and the camera calibration data generated at step S324.

At step S356, head tracker 650 matches the rendered image data generated at step S352 and the camera image data extracted at step S354 to find the camera image data which best matches the rendered head model.

At step S358, head tracker 650 uses the camera image data identified at step S356 which best matches the rendered head model to determine the 3D position and orientation of the participant's head for the current frame of video data.

At the same time that step S358 is performed, at step S360, the positions of the head features in the camera image data determined at step S356 are input into a conventional Kalman filter to generate an estimate of the 3D position and orientation of the participant's head for the next frame of video data. Steps S350 to S360 are performed repeatedly for the participant as frames of

video data are received from video camera 602.

Referring again to Figure 23, at step S336, central controller 636 determines whether there is another participant in the meeting, and steps S326 to S336 are repeated until processing has been performed for each participant in the manner described above. However, while these steps are performed for each participant, at step S334, head tracker 650 continues to track the head of each participant who has already sat down.

When it is determined at step S336 that there are no further participants in the meeting and that accordingly the head of each participant is being tracked by head tracker 650, then, at step S338, central controller 636 causes an audible signal to be output from processing apparatus 624 to indicate that the meeting between the participants can begin.

Figure 26 shows the processing operations performed by processing apparatus 624 as the meeting between the participants takes place.

Referring to Figure 26, at step S370, head tracker 650 continues to track the head of each participant in the

meeting. The processing performed by head tracker 650 at step S370 is the same as that described above with respect to step S334, and accordingly will not be described again here.

5

At the same time that head tracker 650 is tracking the head of each participant at step S370, at step S372 processing is performed to generate and store data in meeting archive database 660.

10

Figure 27 shows the processing operations performed at step S372.

15

Referring to Figure 27, at step S380, archive processor 658 generates a so-called "viewing parameter" for each participant defining at whom the participant is looking.

20

Figure 28 shows the processing operations performed at step S380.

25

Referring to Figure 28, at step S410, archive processor 658 reads the current three-dimensional position of each participant's head from head tracker 650, this being the position generated in the

processing performed by head tracker 650 at step S358 (Figure 25).

At step S412, archive processor 658 reads the current orientation of the head of the next participant (this being the first participant the first time step S412 is performed) from head tracker 650. The orientation read at step S412 is the orientation generated in the processing performed by head tracker 650 at step S358 (Figure 25).

At step S414, archive processor 658 determines the angle between a ray defining where the participant is looking (a so-called "viewing ray") and each notional line which connects the head of the participant with the centre of the head of another participant.

More particularly, referring to Figures 29 and 30, an example of the processing performed at step S414 is illustrated for one of the participants, namely participant 610 in Figure 21. Referring to Figure 29, the orientation of the participant's head read at step S412 defines a viewing ray 690 from a point between the centre of the participant's eyes which is perpendicular to the participant's head. Similarly, referring to

Figure 30, the positions of all of the participant's heads read at step S410 define notional lines 692, 694, 696 from the point between the centre of the eyes of participant 610 to the centre of the heads of each of the other participants 606, 608, 612. At step S414, archive processor 658 determines the angles 698, 700, 702 between the viewing ray 690 and each of the notional lines 692, 694, 696.

Referring again to Figure 28, at step S416, archive processor 658 selects the angle 698, 700 or 702 which has the smallest value. Thus, referring to the example shown in Figure 30, the angle 700 would be selected.

At step S418, archive processor 658 determines whether the selected angle has a value less than 10° .

If it is determined at step S418 that the angle is less than 10° , then, at step S420, archive processor 658 sets the viewing parameter for the participant to the number (allocated at step S304 in Figure 23) of the participant connected by the notional line which makes the smallest angle with the viewing ray. Thus, referring to the example shown in Figure 30, if angle 700 is less than 10° , then the viewing parameter would be set to the

participant number of participant 606 since angle 700 is the angle between viewing ray 690 and notional line 694 which connects participant 610 to participant 606.

5 On the other hand, if it is determined at step S418 that the smallest angle is not less than 10° , then, at step S422, archive processor 658 sets the value of the viewing parameter for the participant to "0". This indicates that the participant is determined to be looking at none
10 of the other participants since the viewing ray 690 is not close enough to any of the notional lines 692, 694, 696. Such a situation could arise, for example, if the participant was looking at notes or some other object in the meeting room.

15 At step S424, archive processor 658 determines whether there is another participant in the meeting, and steps S412 to S424 are repeated until the processing described above has been carried out for each of the participants.

20 Referring again to Figure 27, at step S382, central controller 636 and voice recognition processor 654 determine whether any speech data has been received from the microphones 604 for the current frame of video data.

If it is determined at step S382 that speech data has been received, then, at step S384, archive processor 658 processes the viewing parameters generated at step S380 to determine which of the participants in the meeting is speaking.

Figure 31 shows the processing operations performed at step S384 by archive processor 658.

Referring to Figure 31, at step S440, the number of occurrences of each viewing parameter value generated at step S380 is determined, and at step S442, the viewing parameter value with the highest number of occurrences is selected. More particularly, the processing performed at step S380 in Figure 27 will generate one viewing parameter value for the current frame of video data for each participant in the meeting (thus, in the example shown in Figure 21, four values would be generated). Each viewing parameter will have a value which corresponds to the participant number of one of the other participants or "0". Accordingly, at step S440 and S442, archive processor 658 determines which of the viewing parameter values generated at step S380 occurs the highest number of times for the current frame of video data.

At step S444, it is determined whether the viewing parameter with the highest number of occurrences has a value of "0" and, if it has, at step S446, the viewing parameter value with the next highest number of occurrences is selected. On the other hand, if it is determined at step S444 that the selected value is not "0", then step S446 is omitted.

At step S448, the participant defined by the selected viewing parameter value (that is, the value selected at step S442 or, if this value is "0" the value selected at step S446) is identified as the participant who is speaking, since the majority of participants in the meeting will be looking at the speaking participant.

Referring again to Figure 27, at step S386, archive processor 658 stores the viewing parameter value for the speaking participant, that is the viewing parameter value generated at step S380 defining at whom the speaking participant is looking, for subsequent analysis, for example in memory 642.

At step S388, archive processor 658 informs voice recognition processor 654 of the identity of the speaking participant determined at step S384. In response, voice

5

10

15

20

Referring to Figure 32, meeting archive database 660 stores time information represented by the horizontal axis 800, on which each unit represents a predetermined amount of time, for example one frame of video data received from camera 602. The MPEG 2 data generated at step S390 is stored as data 802 in meeting archive database 660, together with timing information (this timing information being schematically represented in Figure 32 by the position of the MPEG 2 data 802 along the horizontal axis 800).

Referring again to Figure 27, at step S392, archive processor 658 stores any text data generated by voice recognition processor 654 at step S388 for the current frame in meeting archive database 660 (indicated at 804 in Figure 32). More particularly, the text data is stored with a link to the corresponding MPEG 2 data, this link being represented in Figure 32 by the text data being stored in the same vertical column as the MPEG 2 data. As will be appreciated, there will not be any text data for storage from participants who are not speaking. In the example shown in Figure 32, text is stored for the first ten time slots for participant 1 (indicated at 806), for the twelfth to twentieth time slots for participant 3 (indicated at 808), and for the twenty-

first time slot for participant 4 (indicated at 810). No text is stored for participant 2 since, in this example, participant 2 did not speak during the time slots shown in Figure 32.

5

At step S394, archive processor 658 stores the viewing parameter value generated for each participant at step S380 in the meeting archive database 660 (indicated at 812 in Figure 32). Referring to Figure 32, a viewing parameter value is stored for each participant together with a link to the associated MPEG 2 data 802 and the associated text data 804 (this link indicated in Figure 32 by the viewing parameters values being stored in the same column as the associated MPEG 2 data 802 and associated text data 804). Thus, referring to the first time slot by way of example, the viewing parameter value for participant 1 is "3", indicating that participant 1 is looking at participant 3, the viewing parameter value for participant 2 is "1", indicating that participant 2 is looking at participant 1, the viewing parameter value for participant 3 is also "1", indicating that participant 3 is also looking at participant 1, and the viewing parameter value for participant 4 is "0", indicating that participant 4 is not looking at any of the other participants (in the example shown in

10

15

20

25

00532533-032200

Figure 21, the participant indicated at 612 is looking at her notes rather than any of the other participants).

At step S396, central controller 636 and archive processor 658 determine whether one of the participants in the meeting has stopped speaking. In this embodiment, this check is performed by examining the text data 804 to determine whether text data for a given participant was present for the previous time slot, but is not present for the current time slot. If this condition is satisfied for a participant (that is, a participant has stopped speaking), then, at step S398, archive processor 658 processes the viewing parameter values for the participant who has stopped speaking previously stored when step S386 was performed (these viewing parameter values defining at whom the participant was looking during the period of speech which has now stopped) to generate data defining a viewing histogram. More particularly, the viewing parameter values for the period in which the participant was speaking are processed to generate data defining the percentage of time during that period that the speaking participant was looking at each of the other participants.

Figures 33A and 33B show the viewing histograms

corresponding to the periods of text 806 and 808 respectively in Figure 32.

Referring to Figure 32 and Figure 33A, during the period 806 when participant 1 was speaking, he was looking at participant 3 for six of the ten time slots (that is, 60% of the total length of the period for which he was talking), which is indicated at 900 in Figure 33A, and at participant 4 for four of the ten time slots (that is, 40% of the time), which is indicated at 910 in Figure 33A.

Similarly, referring to Figure 32 and Figure 33B, during the period 808, participant 3 was looking at participant 1 for approximately 45% of the time, which is indicated at 920 in Figure 33B, at participant 4 for approximately 33% of the time, indicated at 930 in Figure 33B, and at participant 2 for approximately 22% of the time, which is indicated at 940 in Figure 33B.

Referring again to Figure 27, at step S400, the viewing histogram generated at step S398 is stored in the meeting archive database 660 linked to the associated period of text for which it was generated. Referring to Figure 32, the stored viewing histograms are indicated at 814, with

the data defining the histogram for the text period 806 indicated at 816, and the data defining the histogram for the text period 808 indicated at 818. In Figure 32, the link between the viewing histogram and the associated text is represented by the viewing histogram being stored in the same columns as the text data.

On the other hand, if it is determined at step S396 that, for the current time period, one of the participants has not stopped speaking, then steps S398 and S400 are omitted.

At step S402, central controller 636 determines whether another frame of video data has been received from camera 602. Steps S380 to S402 are repeatedly performed while image data is received from camera 602.

When data is stored in meeting archive database 660, then the meeting archive database 660 may be interrogated to retrieve data relating to the meeting.

Figure 34 shows the processing operations performed to search the meeting archive database 660 to retrieve data relating to each part of the meeting which satisfies search criteria specified by a user.

Referring to Figure 34, at step S500, central controller 636 causes display processor 664 to display a message on display device 626 requesting the user to enter information defining the search of meeting archive database 660 which is required. More particularly, in this embodiment, central controller 636 causes the display shown in Figure 35A to appear on display device 626.

Referring to Figure 35A, the user is requested to enter information defining the part or parts of the meeting which he wishes to find in the meeting archive database 660. More particularly, in this embodiment, the user is requested to enter information 1000 defining a participant who was talking, information 1010 comprising one or more key words which were said by the participant identified in information 1000, and information 1020 defining the participant to whom the participant identified in information 1000 was talking. In addition, the user is able to enter time information defining a portion or portions of the meeting for which the search is to be carried out. More particularly, the user can enter information 1030 defining a time in the meeting beyond which the search should be discontinued (that is, the period of the meeting before the specified time

should be searched), information 1040 defining a time in the meeting after which the search should be carried out, and information 1050 and 1060 defining a start time and end time respectively between which the search is to be carried out. In this embodiment, information 1030, 1040, 1050 and 1060 may be entered either by specifying a time in absolute terms, for example in minutes, or in relative terms by entering a decimal value which indicates a proportion of the total meeting time. For example, entering the value 0.25 as information 1030 would restrict the search to the first quarter of the meeting.

In this embodiment, the user is not required to enter all of the information 1000, 1010 and 1020 for one search, and instead may omit one or two pieces of this information. If the user enters all of the information 1000, 1010 and 1020, then the search will be carried out to identify each part of the meeting in which the participant identified in information 1000 was talking to the participant identified in information 1020 and spoke the key words defined in information 1010. On the other hand, if information 1010 is omitted, then a search will be carried out to identify each part of the meeting in which the participant defined in information 1000 was talking to the participant defined in information 1020

irrespective of what was said. If information 1010 and 1020 is omitted, then a search is carried out to identify each part of the meeting in which the participant defined in information 1000 was talking, irrespective of what was said and to whom. If information 1000 is omitted, then a search is carried out to identify each part of the meeting in which any of the participants spoke the key words defined in information 1010 to the participant defined in information 1020. If information 1000 and 1010 is omitted, then a search is carried out to identify each part of the meeting in which any of the participants spoke to the participant defined in information 1020. If information 1020 is omitted, then a search is carried out to identify each part of the meeting in which the participant defined in information 1000 spoke the key words defined in information 1010, irrespective of to whom the key word was spoken. Similarly, if information 1000 and 1020 is omitted, then a search is carried out to identify each part of the meeting in which the key words identified in information 1010 were spoken, irrespective of who said the key words and to whom.

In addition, the user may enter all of the time information 1030, 1040, 1050 and 1060 or may omit one or more pieces of this information.

5

10

15

20

25

speaking participant looked at other participants while

speaking, provided that the speaking participant looked at the participant defined in information 1020 for at least 25% of the time of the speech. Thus, a period of speech in which the value of the viewing histogram is equal to or above 25% for two or more participants would be identified if any of these participants were specified in information 1020. If the information 1010 has been input by the user, then central controller 636 and text searcher 662 search each portion of text previously identified on the basis of information 1000 and 1020 (or all portions of text if information 1000 and 1020 was not entered) to identify each portion containing the key word(s) identified in information 1010. If any time information has been entered by the user, then the searches described above are restricted to the meeting times defined by those limits.

At step S504, central controller 636 causes display processor 664 to display a list of relevant speeches identified during the search to the user on display device 626. More particularly, central controller 636 causes information such as that shown in Figure 35B to be displayed to the user. Referring to Figure 35B, a list is produced of each speech which satisfies the search parameters, and information is displayed defining the

start time for the speech both in absolute terms and as a proportion of the full meeting time. The user is then able to select one of the speeches for playback by clicking on the required speech in the list using the mouse 630.

At step S506, central controller 636 reads the selection made by the user at step S504, and plays back the stored MPEG 2 data 802 for the relevant part of the meeting from meeting archive database 660. More particularly, central controller 636 and display processor 664 decode the MPEG 2 data 802 and output the image data and sound via display device 626.

At step S508, central controller 636 determines whether the user wishes to cease interrogating the meeting archive database 660 and, if not, steps S500 to S508 are repeated.

Various modifications and changes can be made to the embodiment of the invention described above.

For example, in the embodiment above the microphones 604 are provided on the meeting room table. However, instead, a microphone on video camera 602 may be used to

record sound data.

In the embodiment above, image data is processed from a single video camera 602. However, to improve the accuracy with which the head of each participant is tracked, video data from a plurality of video cameras may be processed. For example, image data from a plurality of cameras may be processed as in steps S350 to S356 of Figure 25 and the resulting data from all of the cameras input to a Kalman filter at step S360 in a conventional manner to generate a more accurate estimate of the position and orientation of each participant's head in the next frame of video data from each camera. If multiple cameras are used, then the MPEG 2 data 802 stored in meeting archive database 660 may comprise the video data from all of the cameras and, at steps S504 and S506 in Figure 34 image data from a camera selected by the user may be replayed.

In the embodiment above, the viewing parameter for a given participant defines at which other participant the participant is looking. However, the viewing parameter may also be used to define at which object the participant is looking, for example a display board, projector screen etc. Thus, when interrogating the

meeting archive database 660, information 1020 in Figure 35A could be used to specify at whom or at what the participant was looking when he was talking.

5 In the embodiment above, at step S502 (Figure 34), the viewing histogram for a particular portion of text is considered and it is determined that the participant was talking to a further participant if the percentage of gaze time for the further participant in the viewing
10 histogram is equal to or above a predetermined threshold. Instead, however, rather than using a threshold, the participant to whom the speaking participant was looking during the period of text may be defined to be the participant having the highest percentage gaze value in
15 the viewing histogram (for example participant 3 in Figure 33A, and participant 1 in Figure 33B).

In the embodiment above, the MPEG 2 data 802, the text data 804, the viewing parameters 812 and the viewing
20 histograms 814 are stored in meeting archive database 660 in real time as data is received from camera 602 and microphones 604. However, instead, the video and sound data may be stored and data 802, 804, 812 and 814 generated and stored in meeting archive database 660 in
25 non-real-time.

In the embodiment above, the MPEG 2 data 802, the text data 804, the viewing parameters 812 and the viewing histograms 814 are generated and stored in the meeting archive database 660 before the database is interrogated to retrieve data for a defined part of the meeting. However, some, or all, of the data 804, 812 and 814 may be generated in response to a search of the meeting archive database 660 being requested by the user by processing the stored MPEG 2 data 802, rather than being generated and stored prior to such a request. For example, although in the embodiment above the viewing histograms 814 are calculated and stored in real-time at steps S398 and S400 (Figure 27), these histograms could be calculated in response to a search request being input by the user.

In the embodiment above, text data 804 is stored in meeting archive database 660. Instead, audio data may be stored in the meeting archive database 660 instead of the text data 804. The stored audio data would then either itself be searched for key words using voice recognition processing or converted to text using voice recognition processing and the text search using a conventional text searcher.

In the embodiment above, processing apparatus 624 includes functional components for receiving and generating data to be archived (for example, central controller 636, head tracker 650, head model store 652, voice recognition processor 654, speech recognition parameter store 656 and archive processor 658), functional components for storing the archive data (for example meeting archive database 660), and also functional components for searching the database and retrieving information therefrom (for example central controller 636 and text searcher 662). However, these functional components may be provided in separate apparatus. For example, one or more apparatus for generating data to be archived, and one or more apparatus for database searching may be connected to one or more databases via a network, such as the Internet.

Also, referring to Figure 36, video and sound data from one or more meetings 1100, 1110, 1120 may be input to a data processing and database storage apparatus 1130 (which comprises functional components to generate and store the archive data), and one or more database interrogation apparatus 1140, 1150 may be connected to the data processing and database storage apparatus 1130 for interrogating the database to retrieve information

therefrom.

In the embodiment above, processing is performed by a computer using processing routines defined by programming instructions. However, some, or all, of the processing could be performed using hardware.

Although the embodiment above is described with respect to a meeting taking place between a number of participants, the invention is not limited to this application, and, instead, can be used for other applications, such as to process image and sound data on a film set etc.

Different combinations of the above modifications are, of course, possible and other changes and modifications can be made without departing from the spirit and scope of the invention.

The contents of the assignee's co-pending PCT application PCT/GB00/00718 filed on 01 March 2000 and designating, inter alia, the United States of America as a designated state, and the contents of the assignee's co-pending US application 09/519,178 (attorney reference number 2635650; CFP 1194 US) filed on 06 March 2000 are hereby

CFP 1240 US (2647350)

incorporated by reference.

[illegible]